

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A distributed multiprocessing computer system, comprising:

a plurality of microprocessor units coupled to each other, wherein each microprocessor unit comprises;

a router to route message packets between said microprocessor units, ~~and wherein said router prioritizes message packets based upon type of message packet, age of the message packet, size of the message packet and source of the message packet;~~

a plurality of network input ports and network output ports connecting said plurality of microprocessor units to form a computer network, wherein each of said network input ports couples to one or more associated local arbiters in the router, each of said local arbiters operable to select a message packet among message packets waiting at the network input port,

wherein the router gives higher priority to message packets associated with the network ports than to message packets associated with other microprocessor ports and wherein the router arbitrates between the network ports using a predetermined prioritization and a least-recently-granted prioritization.

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

2. (Previously presented) The distributed multiprocessing computer system of claim 1 wherein said router includes a plurality of timers that indicate when a message packet must be immediately dispatched.

3. (Previously presented) The distributed multiprocessing computer system of claim 1 wherein said microprocessor unit further includes a plurality of microprocessor input ports and microprocessor output ports that allow the exchange of message packets between hardware functional units in the microprocessor unit and between microprocessor units.

4. (Original) The distributed multiprocessing computer system of claim 3 wherein each of said microprocessor input ports couples to local arbiters in the router, each of said local arbiters able to select a message packet among message packets waiting at the microprocessor input port.

5. (Original) The distributed multiprocessing computer system of claim 4 wherein each of said network output ports and microprocessor output ports couples to a global arbiter in the router that selects a message packet from message packets nominated by the local arbiters of said network input ports and microprocessor input ports.

6. (Original) The computer system of claim 5 wherein if a first message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the first message packet type from the global arbiter of the destination network output port or microprocessor output port.

7. (Original) The computer system of claim 6 wherein if a second message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the second

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

message packet type from the global arbiter of the destination network output port or microprocessor output port.

8. (Original) The computer system of claim 7 wherein if a third message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the third message packet type from the global arbiter of the destination network output port or microprocessor output port.

9. (Original) The computer system of claim 8 wherein if a fourth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the fourth message packet type from the global arbiter of the destination network output port or microprocessor output port.

10. (Original) The computer system of claim 9 wherein if a fifth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the fifth message packet type from the global arbiter of the destination network output port or microprocessor output port.

11. (Original) The computer system of claim 10 wherein if a sixth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the sixth message packet type from the global arbiter of the destination network output port or microprocessor output port.

12. (Original) The computer system of claim 11 wherein if a seventh message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the seventh

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

message packet type from the global arbiter of the destination network output port or microprocessor output port.

13. (Original) The computer system of claim 5 wherein said network output port global arbiter or microprocessor output port global arbiter selects said message packet Least-Recently-Granted from the network input ports, then Least-Recently-Granted from the microprocessor input ports if said network output port or microprocessor output port is idle.

14. (Previously presented) A method of routing messages in a distributed multiprocessing computer system, comprising:

selecting a message packet at each of a plurality of microprocessor router input ports from message packets buffered at each input port based on the type of message packet; and

transmitting from an idle microprocessor router output port a message packet chosen from the plurality of selected microprocessor router input port message packets, said message packet chosen for transmission by the output port based on the microprocessor router input port priority,

wherein said selecting and said transmitting reduces routing latency of the distributed multiprocessing computer system by assigning message packets already in transit higher priority than new message packets and by implementing a combination of type-based and least-recently-granted routing algorithms for both message packets already in transit and new message packets.

15. (Original) The method of claim 14 wherein said selecting a message packet includes the step of:

determining if a Block Response packet is ready to be dispatched from the input port buffer; and

Appl. No. 09/652,322
Amtd. dated September 27, 2004
Reply to Office action of July 22, 2004

if the Block Response packet is ready, selecting the Block Response packet.

16. (Original) The method of claim 15 wherein said selecting a message packet includes the step of:

if no Block Response packet is ready, determining if a Acknowledgment packet is ready to be dispatched from the input port buffer; and
if the Acknowledgment packet is ready, selecting the Acknowledgment packet.

17. (Original) The method of claim 16 wherein said selecting a message packet includes the step of:

if no Acknowledgment packet is ready, determining if an Invalidation Broadcast packet is ready to be dispatched from the input port buffer; and
if the Invalidation Broadcast packet is ready, selecting the Invalidation Broadcast packet.

18. (Original) The method of claim 17 wherein said selecting a message packet includes the step of:

if no Invalidation Broadcast packet is ready, determining if a Forward packet is ready to be dispatched from the input port buffer; and
if the Forward packet is ready, selecting the Forward packet.

19. (Original) The method of claim 18 wherein said selecting a message packet includes the step of:

if no Forward packet is ready, determining if a Request packet is ready to be dispatched from the input port buffer; and
if the Request packet is ready, selecting the Request packet.

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

20. (Original) The method of claim 19 wherein said selecting a message packet includes the step of:

if no Request packet is ready, determining if a Write I/O packet is ready to be dispatched from the input port buffer; and
if the Write I/O packet is ready, selecting the Write I/O packet.

21. (Original) The method of claim 20 wherein said selecting a message packet includes the step of:

if no Write I/O packet is ready, determining if a Read I/O packet type is ready to be dispatched from the input port buffer; and
if the Read I/O packet is ready, selecting the Read I/O packet.

22. (Original) The method of claim 14 wherein said transmitting a message packet includes the step of prioritizing said message packet Least-Recently-Granted from network input ports, then Least-Recently-Granted from microprocessor input ports, wherein said network input ports and said microprocessor input ports are microprocessor router input ports.

23. (Previously presented) A distributed multiprocessing computer system, comprising:

means for selecting a message packet at each of a plurality of microprocessor router input ports from message packets buffered at each input port based on the type of message packet wherein message packets that pass through at least three microprocessors to reach a destination are given highest priority; and

means for transmitting from an idle microprocessor router output port a message packet chosen from the plurality of selected microprocessor router input port message packets, said message packet chosen for transmission by the output port based on the microprocessor router input port priority,

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

wherein operation of said means for selecting and said means for transmitting reduces routing latency of the distributed multiprocessing computer system.

24. (Previously presented) A distributed multiprocessing computer system, with a plurality of microprocessors, which comprise:

- a router to route message packets between said microprocessors, and wherein said router prioritizes message packets based upon type of message packet, age of the message packet, and source of the message packet;

- a plurality of network input ports and network output ports connecting said plurality of microprocessors to form a computer network, wherein each of said network input ports couples to one or more associated local arbiters in the router, each of said local arbiters operable to select a message packet among message packets waiting at the network input port;

- a plurality of microprocessor input ports and microprocessor output ports that allow the exchange of message packets between hardware functional units in the microprocessor and between microprocessors,

wherein each network input port and each microprocessor input port is assigned a different priority and wherein all network input ports have higher priority than the microprocessor input ports; and

a disk drive coupled to each of said plurality of microprocessors.

25. (Previously presented) The distributed multiprocessing computer system of claim 23 wherein the message packets are selected in an order based on the type of message packet and wherein the order from highest priority to lowest priority consists of Block Response packet type, Acknowledgement packet type, Invalidation Broadcast packet type, Forward packet type, Request packet type, Write I/O packet type, and Read I/O packet type.

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

26. (Previously presented) The distributed multiprocessing computer system of claim 24 wherein said router prioritizes message packets based on the source of the message packet comprises the router giving message packets associated with network input ports higher priority than message packets associated with microprocessor input ports.

27. (Previously presented) The distributed multiprocessing computer system of claim 26 wherein said router prioritizes message packets based on the source of the message packet further comprises the router giving message packets associated with I/O ports lowest priority and wherein the I/O ports are associated the disk drive coupled to each of said plurality of microprocessors.

28. (Previously presented) The distributed multiprocessing computer system of claim 2 wherein the plurality of timers comprise starvation timers and drain timers.

29. (Previously presented) The distributed multiprocessing computer system of claim 5 wherein the global arbiter selects and outputs the message packet from message packets based on an input port hierarchy.

30. (Previously presented) A system, comprising:
a plurality of microprocessors coupled to each other and configured to
prioritize message packets transferred between two or more of the
microprocessors and message packets transferred between two or
more hardware units within each of the microprocessors,
wherein each of the microprocessors comprises,
network ports configured to receive message packets from other
microprocessors and output message packets to other
microprocessors;

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

microprocessor ports configured to receive message packets from the hardware units and output message packets to the hardware units; and
a router coupled to the network ports and the microprocessor ports, wherein the router assigns the network ports higher priority than the microprocessor ports and prioritizes message packets associated with each of the network ports and each of the microprocessor ports according to a predetermined prioritization and according to a least-recently-granted prioritization.

31. (Previously presented) The system of claim 30 wherein the predetermined prioritization is based on message packet types.

32. (Previously presented) The system of claim 31 wherein the message packet types are selected from a group of packets types consisting of: block response packets, acknowledgement packets, invalidation broadcast packets, forward packets, request packets, write I/O packets and read I/O packets.

33. (Previously presented) The system of claim 32 wherein block response packets have highest priority.

34. (Previously presented) The system of claim 31 wherein the predetermined prioritization is based on a source of the message packet.

35. (Previously presented) The system of claim 34 wherein the source of the message packet is selected from a group of sources consisting of: a first network port, a second network port, a third network port, a fourth network port, a cache control unit, a memory controller and an input/output port.

36. (Previously presented) The system of claim 35 wherein the first network port, the second network port, the third network port, the fourth network port, the

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

cache control unit, the memory controller and the input/output port are each assigned a different predetermined priority and wherein the first, second, third and fourth network ports are assigned a higher priority than the cache control unit, the memory controller and the input/output port.

37. (Previously presented) The system of claim 36 wherein the cache control unit and the memory controller are assigned a higher priority than the input/output port.

38. (Previously presented) The system of claim 30 wherein the router further comprises:

local arbiters;

global arbiters; and

input buffers,

wherein the local and global arbiters implement algorithms to carry out the least-recently-granted prioritization.

39. (Previously presented) The system of claim 38 wherein the algorithms detect when a message packet is denied transmission for more than a threshold amount of time due to at least one of a type of the message packet and a source of the message packet.

40. (Previously presented) The system of claim 38 wherein the algorithms detect when a message packet is denied transmission for more than a threshold amount of time due to the local arbiters and the global arbiters operating independently of each other.

41. (Previously presented) The system of claim 38 wherein the router is configured to periodically change the value of a bit associated with a message packet to identify the message packet as non-current.

Appl. No. 09/652,322
Amdt. dated September 27, 2004
Reply to Office action of July 22, 2004

42. (Previously presented) The system of claim 41 wherein the router further comprises a starvation timer that counts down from a maximum value to zero after a message packet is identified as non-current.

43. (Previously presented) The system of claim 42 wherein the router further comprises buffers reserved to receive message packets that cause the starvation timer to reach zero.

44. (Previously presented) The system of claim 43 wherein the router further comprises a drain timer that is activated when the starvation timer reaches zero and when a target buffer of a message packet associated with an expired starvation timer is available.

45. (Previously presented) The system of claim 44 wherein, when the drain timer expires, the router dispatches only message packets that are non-current until there are no message packets associated with an expired starvation timer.

46. (Previously presented) The system of claim 38 wherein the local and global arbiters are configured to perform tests before transmitting a particular message packet, wherein the tests are selected from a group of tests consisting of: checking if a network can accept the particular message packet, checking whether the particular message packet is selected by more than one local arbiter, checking a message packet type, checking if an output port of the router is busy, checking if a destination can accept the particular message packet and checking if a starved message packet is waiting at the input buffers.

47. (New) The distributed multiprocessing computer system of claim 1 wherein the router arbitrates between the other microprocessor ports using a predetermined prioritization and a least-recently-granted prioritization.